

An Approach to Security Requirements Engineering for a High Assurance Systemⁱ

Cynthia E. Irvine, Timothy Levin, Jeffery D. Wilson, David Shifflett, Barbara Pereira

Department of Computer Science

Naval Postgraduate School

Monterey, California 93943

Abstract

Requirements specifications for high assurance secure systems are rare in the open literature. This paper examines the development of a requirements document for a multilevel secure system that must meet stringent assurance and evaluation requirements. The system is designed to be secure, yet combines popular commercial components with specialized high assurance ones. Functional and non-functional requirements pertinent to security are discussed. A multi-dimensional threat model is presented. The threat model accounts for the developmental and operational phases of system evolution and for each phase accounts for both physical and non-physical threats. We describe our team-based method for developing a requirements document and relate that process to techniques in requirements engineering. The system requirements document presented provides a calibration point for future security requirements engineering techniques intended to meet both functional and assurance goals.

Key words: security, requirements, assurance, threat, specification, engineering

1. Introduction

Sometimes an organization possesses information of such a critical nature that its inappropriate exposure or modification would cause grave damage to the enterprise. Because networked computer systems are now used to store and manage this sort of highly sensitive information, management may seek assurances that those computers and networks cannot be misused in ways that result in lapses of security policy enforcement that would expose the organization to unacceptable risk. Malicious software has long been understood to be a significant threat to information security [1]. Multilevel secure systems are intended to control the sharing of highly sensitive information in the face of malicious software in a manner commensurate with policies relating to information confidentiality and integrity [2].

High assurance computers and networks are intended for the protection of critical information resources in environments that involve access by individuals with a range of authorizations to enterprise information. A high assurance system is an implementation of the Reference Monitor Concept [3], a notion that describes an idealized security policy enforcement mechanism. This ideal has three characteristics. First, it is always invoked to perform its enforcement duties; thus, the resources it manages are only available via the reference monitor. Second, attackers cannot successfully penetrate the mechanism and thus cause it to fail. Finally, it is small enough to ensure that it will, in fact, correctly enforce policy and that it contains no artifices that might be used to neutralize the policy enforcement mechanism. One objective of high assurance security engineering is to address the threat of system subversion [4]. Examples of system subversion abound in commercial software as is evident from web sites devoted

ⁱ The views expressed in this paper are those of the authors and should not be construed to reflect those of their employers or the Department of Defense. This work was supported in part by the MSHN project of the DARPA/ITO Quorum program and by the MYSEA project of the DARPA/ATO CHATS program.

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 2002	2. REPORT TYPE N/A	3. DATES COVERED -		
4. TITLE AND SUBTITLE An Approach to Security Requirements Engineering for a High Assurance System			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Computer Science Naval Postgraduate School Monterey, CA 93943			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF: a. REPORT unclassified			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 21
b. ABSTRACT unclassified				
c. THIS PAGE unclassified				

to *Easter Eggs* [5]. Thompson [6] expanded on the work of Karger and Schell [7] to amply illustrated how subversion can be cleverly implemented to make detection virtually impossible.

A high assurance system is a real instance that approaches the Reference Monitor Concept ideal through a rigorous security engineering process. Its development begins with the capture of the security policy that needs to be enforced and an interpretation of that policy in terms of a computer system. This produces a formal security policy model and subsequent evidence that policy enforcement objectives have been met. In parallel with that formal approach, the engineering team develops a series of specifications that move from high-level requirements to detailed implementation documents and code. The system requirements specification is the most abstract of these specifications. A system requirements specification for a secure system incorporates security consideration in conjunction with all other requirements. The creation of such a requirements document provides a good example for those studying security requirements concerns.

In many cases, the design and development of high assurance secure systems has taken place behind closed doors [8,9,10]. Although these systems have been evaluated under the guidelines of various criteria intended to assess their security functionality and lifecycle assurance, little documentation is available for public scrutiny beyond the final evaluation reports published at the end of the evaluation process. The actual documents produced during system design and development remain cloaked in proprietary secrecy and government constraints and non-disclosure agreements.

Three of the authors of this paper have been members of engineering teams whose previous work resulted in the design and implementation of a proprietary highly trusted secure system [10]. Since entering a less restrictive environment, our collaborators and we have commenced the development of the Multilevel Secure Local Area Network, a high assurance network that relies on commercial-off-the-shelf (COTS) hardware and software to present the user interface. This work has been undertaken as a university research project, under government sponsorship, with the explicit intent of incorporating the commercial/industrial product development practices (e.g., per [10]) that are conducive to the production of highly trusted secure systems. This paper describes the general method used to develop the high-level requirements specifications for a high confidence secure system.

Using our MLS LAN System Requirements Document (see Appendix) for the high assurance multilevel secure local area network as a backdrop and example, we outline the steps needed to move from an intuitive, *ad hoc* notion of the system requirements to a precise, coherent document that can be used for more detailed system specifications and design. The system requirements specification captures not only the system's functional requirements but the non-functional requirements as well.

The construction and use of a system requirements document for a trusted system does not in and of itself ensure that the resulting system is secure, and the point of this paper is not to convince the reader that the target system is secure. In the end, formal and informal analysis (including testing) of the system, its design, its usage documentation and its requirements are necessary to understand whether or not the constructed system is "secure" with respect to its policy. Thus, examination of the requirements specification, and/or the requirements documentation process (e.g., this paper) will not lead to conclusive understanding of the security provided by the described system. A comprehensive formal evaluation of the MLS LAN (e.g., see [11]), which would be the ultimate assessment of the requirements documentation methodology described here, has not been completed. We note that the evidence described in this paragraph provides evaluators with a basis for assessing the assurance or confidence one may ascribe to a security implementation.

The remainder of this paper presents our requirements development method. To put the requirements specification process in the proper context, background information about the MLS LAN is provided in Section 2. Techniques used to create the specification are discussed in Section 3 with reference to Appendix A, an image of the requirements specification document. This section presents our multi-dimensional threat model as well as the method we used to develop the specification. Some general observations about high assurance security requirements specifications and recommendations for process improvement follow in Section 4. Our conclusion is presented in Section 4.2.

2. Overview of the Multilevel Secure LAN

The Center for INFOSEC Studies and Research (CISR) Multilevel Secure Local Area Network (MLS LAN) project [12,13,14,15] at the Naval Postgraduate School is an effort to develop a high assurance LAN that leverages COTS components. The project was driven by two overarching requirements. The first was to protect security-critical information from unauthorized disclosure or modification. The second was to provide users with the ability to use popular COTS office productivity tools for information organization and management at authorized sensitivity levels.

To achieve these objectives, we undertook the design and development of a prototype infrastructure for applications: a local area network that incorporates both high assurance security policy enforcement mechanisms and COTS software and hardware platforms for users' desktops. The former was intended to leverage the considerable government investment in highly trusted systems, while the latter objective resulted from the realization that unless a secure system offered users the same sort of convenient interfaces users have come to expect when handling *normal* information, the secure system would be likely to fail due to lack of *user acceptability* [16].

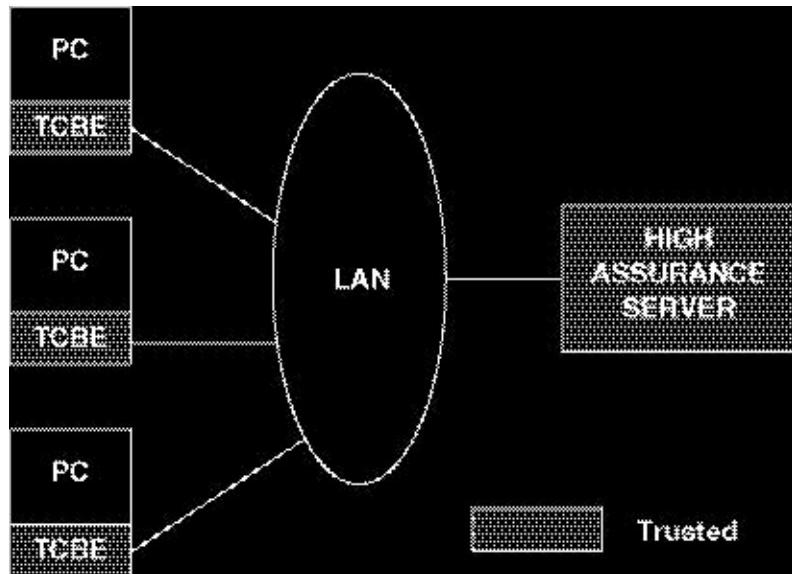


Figure 1: CISR MLS LAN Architecture

The CISR MLS LAN provides controlled sharing of labeled information while permitting users to access that information through popular PC-based COTS personal and office productivity applications. Its architecture is illustrated in Figure 1 [14].

The high assurance server enforces the security policy and controls access to information. It is the core trusted computing base (TCB) for the distributed system. Application protocols run on the High assurance server and provide services and access to shared resources. Each PC is equipped with a trusted computing base extension (TCBE) plug-in board that provides TCB support at the workstation. From these clients, users log on to the TCB, establishing an identity for audit and access control purposes. The TCB components, either the high assurance server base or the TCB extensions, are the only components directly connected to the physical network. Individual components are discussed below.

2.1 Trusted Servers

Each Trusted Server consists of the high assurance TCB, which enforces critical security policy, and untrusted application server instances (viz., one per security level per user) constrained by the TCB. The server supports sharing and labeling and is functionally equivalent in terms of overall application-level protocol support to a COTS application server for the particular protocol it is providing. Thus, it is compatible with existing COTS client packages.

Included in the application servers we have adapted to the high assurance environment are: Internet Mail Access Protocol (IMAP) [17] based on a port of the University of Washington IMAP server [18], Hypertext Transfer Protocol (HTTP) providing an Apache-based port [19], and Simple Mail Transfer Protocol (SMTP) based on sendmail [20]. The servers required little or no code modification to be adapted to the multilevel environment. With a proper configuration, users can view information at or below their current session levels.

2.2 High Assurance Base

The Wang XTS-300 provides our high assurance base [21]. This TCB, by virtue of the protection domains it creates, provides confidence that malicious code will neither cause the exfiltration of sensitive data nor the corruption of information of higher integrity. Thus, one has confidence of correct security policy enforcement.

The high assurance server is defined by the broad properties needed for a viable commercial product. Our definition of a high assurance base is a TCB on the Evaluated Products List (EPL) [22] with a Class B3 or higher digraph based on an evaluation against either the Trusted Computing System Evaluation Criteria (TCSEC) [11] or its network interpretation [23], or equivalent Common Criteria requirements [24].

Minor modifications to XTS-300 TCB networking interfaces contribute to the support of the following desired functions: (1) a trusted path between client workstations and the XTS-300, (2) session-level negotiation at the XTS-300 from the client workstations, and (3) single-level session communications on the Ethernet for client workstations at different session levels. Our modifications permit multiple clients at different access classes to securely communicate with the server through a single physical network device [17]. As stated in Section 1, while these modifications are designed to be secure, their formal verification has yet to be completed [14].

2.3 Client Workstations

Client workstations are typical COTS PCs hosting a popular commercial operating system and a commercial application suite. For mail services the clients include: Lotus Notes, Outlook, Pine, Postal, and Netscapeⁱⁱ. A typical browser supports the client interface to web pages.

To insure that object reuse requirements are met, workstations are considered, in effect, “diskless,” with sufficient volatile RAM-disk capability to support a wide variety of user applications. The workstation TCB extension satisfies object reuse requirements by ensuring that RAM and other volatile primary and secondary storage at the workstation are purged with each change of session level or new user login at the workstation.

2.4 Trusted Computing Base Extension

To extend the TCB across the network, the architecture includes a trusted computing base extension (TCBE) at each COTS workstation [25,15,26]. This component is planned to provide the following services:

- A secure attention key (SAK) that permits users to establish unambiguous communication with the high assurance TCB for unspoofable presentation and capture of security critical data at the user interface.
- Non-bypassable, controlled access to the LAN.
- Protected communication channels between the TCB and the TCBE. These protected communications are based upon protocols that support both the establishment and maintenance of a trusted path and session-level communications.

ⁱⁱ These application names: Lotus Notes, Outlook, and Netscape, are trademarked by their respective owners.

- Mechanisms to ensure high assurance object reuse at the client PC for both primary and secondary storage. While the implementation of these mechanisms is still in progress, experiments by Agacayak [27] indicate that this is possible with the addition of an add-on card.
- Control of the client and its resources at the time of boot, as well as control over security critical actions throughout the client session.
- Return to a secure initial state after an interruption in normal processing (e.g., system crash).

We concluded that a carefully written and reviewed system requirements document is necessary to guide the development of the system architecture and various communications protocols needed in the trusted LAN.

3. Requirements Specification

In this section we provide some insights into our system requirements specification process. The motivation for the process is described. Functional and non-functional requirements are discussed. Our multidimensional threat model is presented. Finally we describe our requirements specification method.

3.1 Motivation and Objectives

A principal motivation for producing a system requirements specification was to provide focus for a complex effort. As we began exploring the design of a group of distributed mechanisms that had to result in coherent enforcement of the network security policy, it became clear that unless each mechanism had the same conceptual system as a goal, the result might be ineffective. In order to create a functional specification from which we could proceed, a requirements document that could be shared among all members of the engineering team was produced.

Another driver for the process was the desire to have a tool that would enable us to judge when we were done building the LAN. The students were particularly concerned about *requirements drift* and *requirements inflation* [28] that might have created a Sisyphean, never-ending, task. The requirements specification also allowed us to examine the system in the abstract to determine whether some essential element had been omitted.

Several of us viewed requirements specification development for high assurance secure systems as a social process within the context of a larger engineering design and development effort and wanted to teach this aspect of high assurance system requirements specification to a new generation. We are not aware of any books or tools that can convey the social and interactive nature of this process as well as actual participation in the effort.

The objective of the requirements specification process was to produce a high level yet rigorous description of system behavior that is complete, consistent, and well formed. Our project objective was to produce a framework for an enabling technology to both support desired user applications and enforce the security policy. The requirements specification describes what we are trying to do, not how we are going to do it.

3.2 Functional and Non-Functional Security Requirements

Some past requirements-research efforts have stated very broadly that all security requirements are *non-functional* [29,30]. On the other hand, it has been suggested that “functions” implement system state changes and that *functional requirements* define how those functions will behave and what states can be reached [31]. In the latter case, certain security requirements **are** functional, as they affect changes in state. Consider for example an attempt to gain access in a system enforcing a mandatory confidentiality policy. Models for implemented systems [32,33] describe system state changes while preserving system security properties. While the state of the system is being functionally transformed, *non-functional requirements* ensure that a set of orthogonal system properties are maintained.

To illustrate the nature of functional and non-functional requirements, consider two systems: a system enforcing a mandatory confidentiality policy (designated SECURE) and a system to automate traffic lights (designated TRAFFIC). Table 1 illustrates a few requirements for each. For the traffic light system, the requirement that the traffic light turns yellow after it has been green for three minutes is a state-changing requirement and is therefore functional. The accuracy of the clock that must be used to measure when three minutes have elapsed presents a non-functional, non-state changing requirement. For the secure system, users must present a valid password in order to

log in to the system. This is a requirement which must be met in order for a state change from *user-logged-on* to *user-logged-in* to occur. Similarly, the requirement that secret information be stored in objects labeled SECRET affects whether information can move from an unstored to a stored state. In contrast, the requirement for the accumulation of an audit trail recording all access attempts does not affect system state. When access attempts succeed, state changes and when they fail state is unchanged. In either case, an audit record is generated.

Table 1: Functional and Non-Functional Requirements

Type	Requirement	System
Functional	A user shall present a valid password in order to log in to the system.	SECURE
	Secret information shall be stored in objects labeled SECRET.	SECURE
	The traffic light shall turn yellow after it has been green for three minutes.	TRAFFIC
Non-Functional	Audit records shall be recorded for all access attempts.	SECURE
	The accuracy of the clock shall be to within five seconds.	TRAFFIC

Thus, we conclude that the construction of a high assurance system intended to enforce national security policy imposes both functional and non-functional requirements on a system. However, our requirements specification is not organized to distinguish functional from non-functional security requirements, as we did not consider that to be an advantage for this project.

3.3 Threat Model

The system requirements specification is motivated by an organization of various classes of threats, characterized as a “threat model.” The threat model is derived through extensive experience with both trusted system design elements and vulnerability analyses of the target environments. The basic elements of the engineering process [28,34] are extended to address the threat model. Figure 2 illustrates our modification to the requirements process.

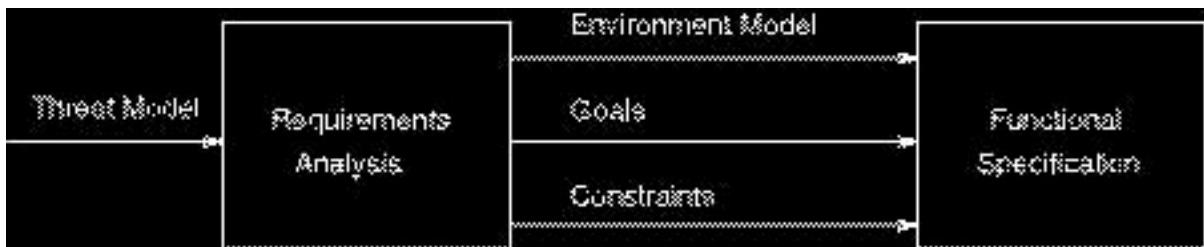


Figure 2: Requirements Process

An understanding of the threat model within the system’s developmental and operational domains drives system requirements. The threat model deals with abstract classes of threats with respect to these domains, such as subversive attacks against the system code base and malicious attacks against the system interface, rather than detailed exploitation mechanisms, such as identity spoofing or elevation of privilege. As stated in Section 2, the objective of the system, and hence the focus of threat analysis (which may deal with more detailed exploitations), is to enforce a policy regarding the disclosure and modification of sensitive information. Availability is not a primary concern, as is reflected in the threat model.

The threats to be addressed are of two kinds, representing the domains of concern. Developmental threats form the first major class of threats to the TCB. These are threats to the ability of the system to fulfill its design requirements. There are many design and implementation mistakes that can occur during a standard software development project [35]. The threat of subversion by insiders [4, 36] is a major concern during the development of secure systems, considering that the objective is to construct a system that is self-protecting [16]. To counter mistakes due to poor engineering, we use a rigorous engineering process. To counter subversion, we use configuration management and quality assurance procedures that require, for example, two-person authorization and other redundant safeguards at key control points. Much like an internal banking system, these safeguards are designed to protect against antagonistic individual insiders, rather than massive collusion.

Development threats are illustrated in Figure 3. Interfaces to the TCB are shown and threats utilizing those interfaces are illustrated in hexagons. These threats include both physical and non-physical attacks. It is worth noting that adequate configuration management and quality assurance can mitigate many physical threats during the system development phase. Additional measures, including procedural security policies, to address physical threats include traditional disaster recovery methods [37] such as: fire and earthquake protection, off-site backups, and locked offices and laboratories.

A comprehensive framework is needed during the development phase of a secure system to counter developmental threats. A requirements specification for this framework is beyond the scope of the example in Appendix A, but some guidance for high assurance development appears in standard evaluation criteria [11,24].

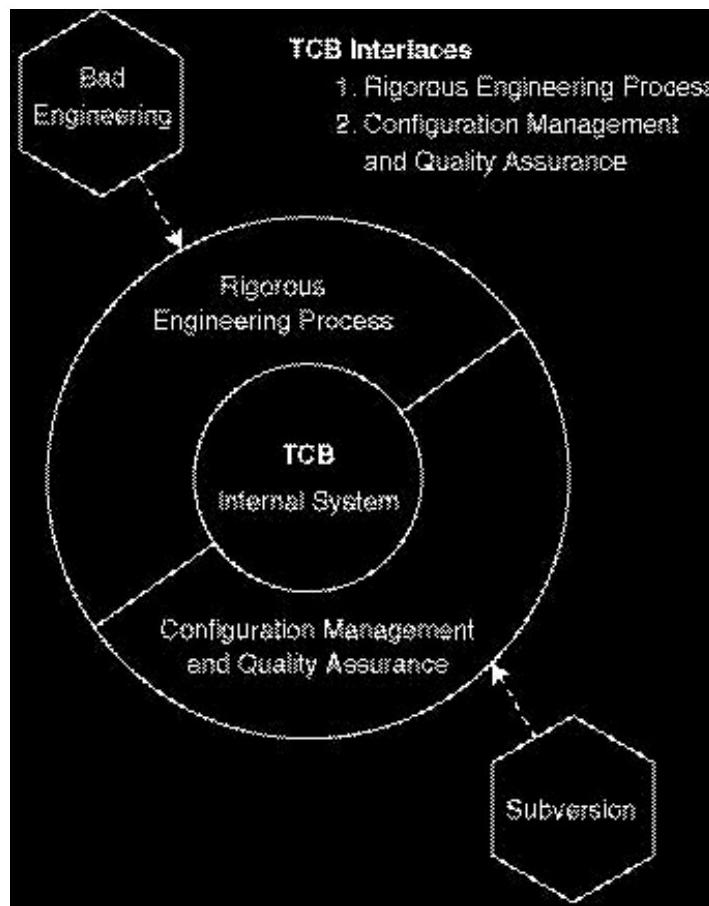


Figure 3: Development Threats

The second, and more obvious, class of threats is operational threats. These are the threats that the system is designed to counter during its operation, and that typically come to mind when one thinks of system security. In the case of the MLS LAN, a high assurance multilevel system, the threats fall into three classes: network threats, malicious software, and user or application misbehavior. Network threats are attacks to the communications protocols within the LAN. For example, an attempt could be made to use a non-TCBE equipped workstation attached to the LAN to modify or collect communications traffic. Malicious software is the principal operational threat to the server. This software would attempt to violate security policies for information confidentiality or integrity by obtaining unauthorized access to information. Trojan Horse software represents a classic example of malicious software and can be used to either directly or indirectly access information [38]. Misbehavior applies to the MLS LAN Workstation where both user actions and malicious software may be used as attack paths. In this case, either users or software attempt to nullify the TCB Extension in order to gain unauthorized access to protected information.

We further decompose the operational threats into software threats and physical threats. Figure 4 illustrates the system of interest, i.e., the trusted computing base (TCB), its interfaces, and the operational software threats to the system by external elements that will drive the requirements specification. Thus the TCB is the internal system. Its interfaces are to the communications protocols, application protocol server, and MLS LAN Workstation. (A direct correspondence between Figure 4 and the second figure of the requirements document (MLS LAN Component Overview) can be observed.) External threats are network attacks, malicious software and misbehavior. The requirements document, when complete, addresses these threats. For example, in the requirements document, the section on MLS LAN Connection Protocol Requirements addresses mitigation of the threat of network attacks on the system.

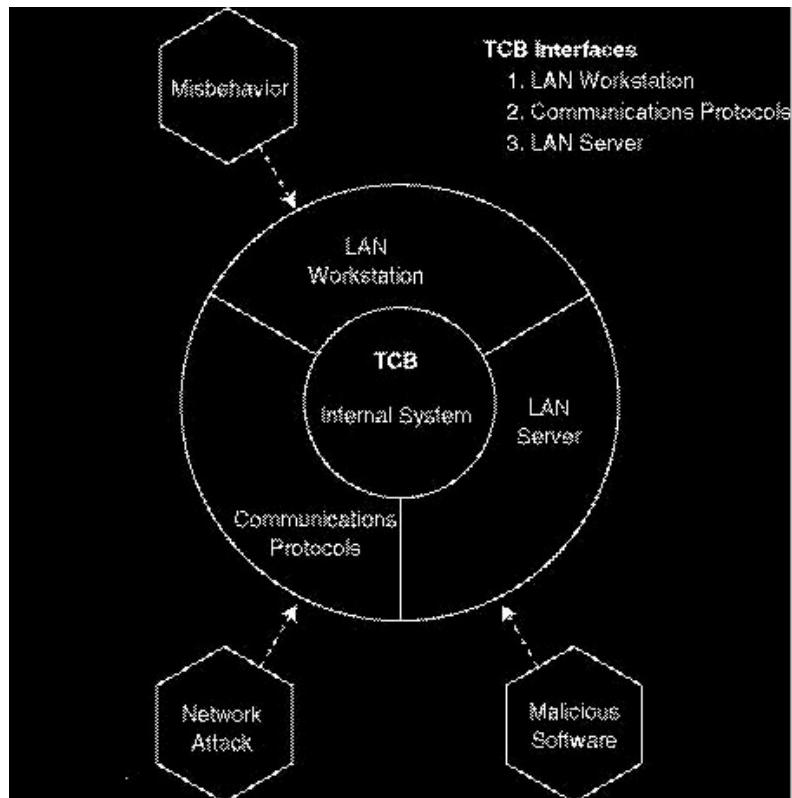


Figure 4: Operational Software Threats

An oft-neglected aspect of system lifecycle management that requires attention is the prospect of physical access to and tampering with the system. For example, if an individual were to break the seals on the XTS-300 or the TCBE and replace the chip containing the BIOS, this could render the system insecure. So, in addition to the usual configuration management concerns the issue of physical protection of the fielded trusted system components must be addressed. Physical threats complement software threats to the operational system, and are illustrated in Figure 5. The TCB and its interfaces are the same, but the external threats are tampering with the communications lines, and TCB components, i.e., the XTS-300 or the TCB Extensions.

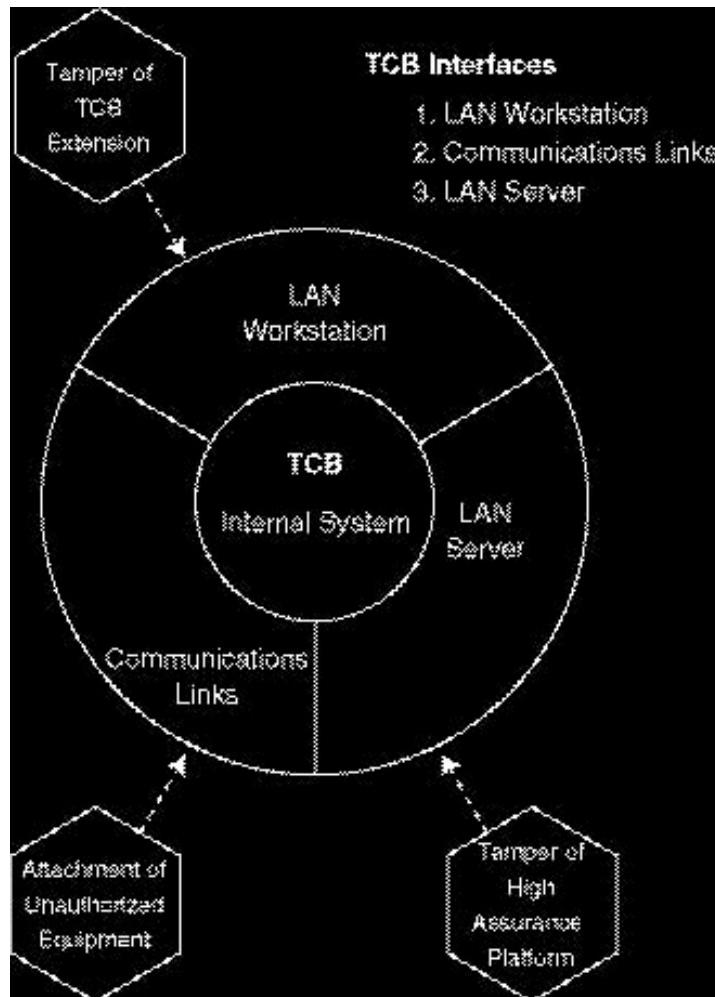


Figure 5: Operational Physical Threats

In summary, our threat model accounts for the developmental and operational stages of system evolution, and for each stage, accounts for both physical and non-physical threats. As the customer objective was not the enforcement of an availability policy, the threat model for this project does not address denial of service concerns.

3.4 Method

Using a requirements specification format [39], we started with a skeletal requirements specification and filled in sections. We often found that requirements from one section would affect decisions associated with another aspect of the system. Thus the requirements specification process was iterative within the component.

We used a two-level “iterative” waterfall process [40] to ensure that the requirements specification was realistic with respect to the design. This hybrid process combines elements of the classic waterfall and the spiral [41] development processes to allow objective-driven stepwise refinement, as well as “upstream” feedback. After completing what appeared to be a reasonable draft of the system requirements specification, we moved to the development of a *functional design specification* for various communications protocols and protection mechanisms of the distributed TCB. (Drafts of both the requirements specification and the functional specification appeared as appendices to a student thesis [42].) This was a much more detailed and concrete design statement that provided insights into the implementation details we were working toward. Those realities fed back into the (upstream) system requirements document in interesting ways.

1. Design decisions gone awry sometimes indicated the need for a guiding principle at the requirements level. For example, the need to repair an overly complex design for storing user data led to the requirement for the Session Database Server to maintain data on each User Session (A.3.2.3.3).
2. Fundamental requirements statements moved out of design and up to requirements. For example, repeated statements about the need for a Secure Attention Key mechanism in different circumstances at the design level were replaced by a single statement at the requirement level (A.3.2.1.1)
3. Conflict with real world possibilities resulted in clarification or refinement at the requirements level. For example, the possibility of session corruption resulting from a broken connection led to the inclusion of a requirement to sever network access to resources if the underlying communication channel is broken (A.3.2.1.3).

Concurrent development of a functional specification allowed us to identify notions that could potentially be generalized in the requirements document. Conversely, items more appropriate for the functional specification were removed from the requirements document. For the less experienced members of the team, the temptation to include implementation details as requirements was enormous. It is fun to think about the implementation, but it is infinitely harder to describe the abstract system. On the other hand, feedback from the more detailed design phase plays a key role. This iterative waterfall approach permitted us to develop documents that would be suitable for evolutionary engineering processes [43,44] as experiential or environmental factors lead to requirements for new versions of the system.

Abstraction is one of the principal objectives of a high-level requirements specification and one of the most difficult to achieve. Although our specification is intentionally abstract it is not intended to be vague, but instead semantically precise. There should be no ambiguity with respect to requirements that the system absolutely had to meet. In addition, the specification should be sufficiently abstract that a variety of implementations could satisfy the requirements [31]. For example, our specification states in Section 3.2.1.1 “the TCB shall provide a Secure Attention Key (SAK) mechanism to invoke a trusted path from workstations to which the TCB has been extended.” Nowhere is there a statement regarding how that SAK should be implemented: that is the purview of the design team. It could be invoked from a keyboard using a combination of keys such as “CTRL-ALT-DEL” or it could be invoked using a special button.

Another aspect of the system that is abstracted away is the functionality associated with the applications. They are relevant to the framework under consideration only in so far as they impose requirements for support from the network in the form of resources for processing capability, memory and network bandwidth.

A key part of the process of developing the requirements specification involved recognition of implementation details. Through a process of winnowing, we were able to avoid inclusion of the implementation details. Some detailed descriptions were transferred to nascent design documents; some discarded; and others were generalized and abstracted to become true requirements. For example, instead of stating what application protocol services were to be supported, the specification states in Section 3.1.3 that “the MLS LAN shall provide...application protocol services ...”.

A fundamental concept for requirements development is that of completeness with respect to the goals, as well as well as functionality. In reviewing progress and draft versions, tests for completeness can include *gedanken* exercises, such as: “can a useless or insecure system be built to this abstract specification?” A mapping of the requirements specification to the threat models is used to ensure that all threats are accounted for. The individual and group reviews ensured that these questions were asked regularly throughout the requirements development process.

3.4.1 Social Process in Requirements Specification

Ultimately, system implementation is the job of the entire team, however during the specification process, it is necessary to consider the perspectives of a larger set of individuals. During the requirements specification process, different members of our specification team were assigned to represent the needs of the various stakeholders in the ultimate product: customers, users, developers, certifiers and accreditors. The customer must be represented because, otherwise the system might become an amusing sandbox for implementation team experimentation and the likelihood that the system would actually be used would be substantially reduced. The engineering team requires representation because this is the group that must be able to move from the abstract specification to a concrete implementation. Certifiers and accreditors must be represented because they are the ones responsible for attesting that the implementation is faithful to the specification. Various sets of system users are represented and include typical users, operators, and administrators. The most senior members of the team, who serve as system architects, represented the government customer. Junior members of the group represented developers whose task would be system implementation. Our experience was that the active representation of all stakeholders helped to ensure that the specification was well balanced with respect to the various, and sometimes conflicting, needs of these stakeholders.

For several months we met on a weekly basis to discuss each new draft of the document. This type of teamwork is best achieved in an environment in which criticism is viewed not as a subjective attack on the writer but as an objective, scientific attempt to achieve the best possible result. As the team matured, criticism was accepted and members with strengths in particular areas were able to contribute in corresponding ways to the emerging document.

The team approach adds to the assurance of the resulting system and mitigates development threats. Because a group of system architects inspected the documents and discussed its semantics, the addition of a subversive artifact during the design stage is considerably more difficult [4]. This threat is directly addressed by way of the project’s configuration management and quality assurance system.

3.4.2 Implicit Requirements

Inspection of the MLS LAN System Requirements Document shows that a number of the functional requirements are quite general. The system must support application protocols, but it is not necessary to specify which protocols. The system must support inexpensive commercial PCs as user workstations. It must support up-to-date versions of commercial operating systems. It follows implicitly that client applications can be up-to-date. The latter two requirements encompass not only a requirement for immediate system performance, but also one for adaptability in that both the COTS operating system and the applications it supports must be upgradeable at any time.

4. Summary

When attempting to build a high assurance system, the objective is to provide a high level of confidence that security policy will be correctly and continuously enforced. Minimization has been recognized as an effective means of ensuring that the system can be judged for correctness and completeness [45]. In a minimized system we ask two questions: first, are all components within the system boundary, i.e., the Trusted Computing Base perimeter, necessary for the correct enforcement of security policy, and, second, are the mechanisms organized such that they are sufficient for security policy enforcement. If additional functionality is added to the secure system, then additional effort will be required to demonstrate that the result meets the assurance requirements. The task in developing a requirements document is to restrict the design of the system by providing sufficient detail so that system specifications are well focused.

4.1 A Calibration Point

Over the past decade there has been significant work in the area of requirements engineering [30] and new efforts are underway to extend these achievements to the area of security requirements engineering. What metric should be applied to judge the effectiveness of these new security requirements engineering techniques? Will those requirements be sufficiently abstract to permit the development of a wide range of system implementations? Will those requirements introduce restrictions that will facilitate the development of system specifications and other documents [31]? The requirements document provided here is a worked example in that it has provided both the abstraction and restrictions necessary to shorten the effort required to specify and construct a high assurance multilevel system. Thus, our specification can provide a calibration point for new techniques in requirements engineering.

4.2 Conclusions

This paper examines the development of a requirements document for a multilevel secure system that must meet stringent assurance requirements. The system is designed to be secure, yet combines popular commercial components with specialized high assurance ones. We describe security objectives as having both functional and non-functional requirements. A multi-dimensional threat model that accounts for developmental and operational phases of system evolution and considers both physical and non-physical threats is presented. Physical threats to the system are explicitly addressed in the requirements engineering process. We describe our team-based approach to system specification and design. By assuming the views of various stakeholders in the system, through open, non-judgmental discourse, and by using the threat model and the high-level design specifications as a check, we have developed an abstract requirements specification. The specification of the NPS MLS LAN provides a worked example of a requirements document for a high assurance secure system and thus may be unique in the open literature. Our specification can be used as a calibration point for future security requirements engineering techniques intended to meet both functional and assurance goals.

5. Acknowledgements

The authors wish to thank Dr. Luqi for useful discussions of the requirements engineering process. The authors are grateful to their U.S. Navy and government sponsors for their support of this research.

References

- [1] W. H. Ware. Security Controls for Computer Systems: Report of Defense Science Board Task Force on Computer Security. Technical Report R-609-1, Rand Corporation, Santa Monica, CA, 1970.
- [2] D. H. Brinkley and R. R. Schell. Concepts and Terminology for Computer Security. In Abrams, Jajodia, and Podell, editors, *Information Security: An Integrated Collection of Essays*, pages 40–97. IEEE Computer Society Press, Los Alamitos, CA, 1995.
- [3] J. P. Anderson. Computer Security Technology Planning Study. Technical Report ESD-TR-73-51, Air Force Electronic Systems Division, Hanscom AFB, Bedford, MA, 1972. (Also available as Vol. I, DITCAD-758206. Vol. II, DITCAD-772806).
- [4] P. D. Myers. Subversion: The Neglected Aspect of Computer Security. Master's thesis, Naval Postgraduate School, Monterey, CA, 1980.
- [5] The Easter Egg Archive. <http://www.eeggs.com/>, last modified 19 May 2000.
- [6] K. Thompson. Reflections on Trusting Trust. Communications of the ACM., 27(8): 761–763, 1984.
- [7] P. A. Karger and R. R. Schell, Multics Security Evolution: Vulnerability Analysis, ESD-TR-74-193, Vol. II, Electronic System Division, Air Force Systems Command, Hanscom AFB, Bedford MA, June 1974
- [8] National Computer Security Center. Final Evaluation Report: Boeing Space and Defense Group, MLS LAN Secure Network Server System, 28 August 1991.
- [9] Gemini Trusted Network Processor (GTNP). In *Information Systems Security Products and Service Catalog Supplement*, Report No.CSC-PB-92/001. April 1992. 4-SUP-3a.3.

- [10] National Computer Security Center. Final Evaluation Report of Gemini Computers, Incorporated Gemini Trusted Network Processor, Version 1.01, 28 June 1995.
- [11] National Computer Security Center. Department of Defense Trusted Computer System Evaluation Criteria, DoD 5200.28-STD, December 1985.
- [12] C. E. Irvine, J. P. Anderson, D. Robb, and J. Hackerson. High Assurance Multilevel Services for Off-The-Shelf Workstation Applications. In Proceedings of the 20th National Information Systems Security Conference, pages 421–431, Crystal City, VA, October 1998.
- [13] J. P. Downey and D. A. Robb. Design of a High Assurance Multilevel Mail Server (HAMMS). Master's thesis, Naval Postgraduate School, Monterey, CA, 1997.
- [14] S. Bryer-Joyner and S. Heller. Secure Local Area Network Services for a High-Assurance Multilevel Network. Master's thesis, Naval Postgraduate School, Monterey, CA, March 1999.
- [15] S. Balmer. Framework for a High-Assurance Security Extension to Commercial Network Clients. Master's thesis, Naval Postgraduate School, Monterey, CA, September 1999.
- [16] J. H. Saltzer and M. D. Schroeder. The Protection of Information in Computer Systems. Proceedings of the IEEE, 63(9): 1278–1308, 1975.
- [17] IMAP Information Center. <http://www.washington.edu/imap/>.
- [18] B. Heads. Developing a High Assurance Multilevel Mail Server. Master's thesis, Naval Postgraduate School, Monterey, CA, March 1999.
- [19] E. Bersack. Implementation of a HTTP (Web) Server on a High Assurance Multilevel Secure Platform. Master's thesis, Naval Postgraduate School, Monterey, CA, December 2000.
- [20] E. Brown. SMTP on a High Assurance Multilevel Server. Master's thesis, Naval Postgraduate School, Monterey, CA, September 2000.
- [21] National Computer Security Center. Final Evaluation Report of HFSI XTS-200, CSC-EPL-92/003 C-Evaluation No. 21-92, 27 May 1992.
- [22] Evaluated Products List, National Computer Security Center. <http://www.radium.ncsc.mil/tpep/epl/>.
- [23] National Computer Security Center. Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria, NCSC-TG-005, July 1987.
- [24] ISO/IEC 15408 - Common Criteria for Information Technology Security Evaluation. Technical Report CCIB-98-026, May 1998.
- [25] J. Hackerson. Design of a Trusted Computing Base Extension for Commercial Off-The-Shelf Workstations (TCBE). Master's thesis, Naval Postgraduate School, Monterey, CA, September 1997.
- [26] B. Turan. Client Bootstrap Under TCBE Control. Master's thesis, Naval Postgraduate School, Monterey, CA, March 2000.
- [27] C. Agacayak. TCBE Control of Object Reuse in Clients. Master's thesis, Naval Postgraduate School, Monterey, CA, March 2000.
- [28] V. Berzins and Luqi. Software Engineering with Abstractions. Addison Wesley, Reading, Massachusetts, 1990.
- [29] J. Mylopoulos, L. Thung, and B. Nixon. Representing Using Nonfunctional Requirements: A Process-Oriented Approach. IEEE Transactions on Software Engineering, 18(6): 483–497, June 1992.
- [30] A. van Lamsweerde. Requirements engineering in the year 00: A research perspective. In Proc. ICSE'2000 - 22nd International Conference on Software Engineering, pages 5–19, Limerick, Ireland, June 2000. ACM Press.
- [31] D. Parnas and J. Maday. Functional documents for computer systems, in Science of Computer Programming, volume 25, October 1995, chapter pages 41–61.
- [32] D. Bell and L. LaPadula. Secure Computer Systems: Mathematical Foundations and Model. Technical Report M74-244, MITRE Corp., Bedford, MA, 1973.
- [33] D. Bell and L. LaPadula. Secure Computer System: Unified Exposition and Multics Interpretation. Technical Report ESD-TR-75-306, MITRE Corp., Hanscom AFB, MA, 1975.
- [34] I. Sommerville. Software Engineering. Addison-Wesley, Reading, MA, Fifth edition, 1995.
- [35] N. Leveson. Safeware. Addison Wesley, Reading, Massachusetts, 1995.
- [36] E. A. Anderson, A Demonstration of the Subversion Threat: Facing a Critical Responsibility in the Defense of Cyberspace, Master's thesis, Naval Postgraduate School, Monterey, CA, March 2002.

- [37] C. Pfleeger. Security in Computing. Prentice Hall, Inc., Englewood Cliffs, NJ, 2nd edition, 1986.
- [38] B. Lampson. A Note on the Confinement Problem. Communications of the A.C.M., 16(10): 613–615, 1973.
- [39] K. Heninger. Specifying Software Requirements for Complex Systems: New Techniques and their Applications. IEEE Transactions on Software Engineering, 2(1): 2–12, January 1980.
- [40] W. Royce. Managing the Development of Large Software Systems: Concepts and Techniques. In Proceedings WESCON, August 1970.
- [41] Boehm, Barry W., "A Spiral Model of Software Development and Enhancement", SIGSOFT Software Engineering Notes, 11, 22-42, 1986.
- [42] J. Wilson. Trusted Networking in a Multilevel Secure Environment. Master's thesis, Naval Postgraduate School, Monterey, CA, June 2000.
- [43] Luqi. Software Evolution Through Rapid Prototyping. IEEE Computer, 22(5): 13–25, May 1989.
- [44] Luqi. A Graph Model for Software Evolution. IEEE Transactions on Software Engineering, 16(8): 917–927, August 1990.
- [45] M. Gasser. Building a Secure Computer System. Van Nostrand Reinhold, New York, NY, 1988.

APPENDIX

A. MLS LAN System Requirements Document - Version 0.2

A.1 Introduction

A.1.1 Purpose

The purpose of this System Requirements Document is to define the design requirements for the Naval Postgraduate School Center for INFOSEC Studies and Research (CISR) Multilevel Secure Local Area Network (MLS LAN) Project.

A.1.2 Scope

This requirements document provides extensive information concerning the design requirements for each of the components of the MLS LAN project. It outlines the mandated system goals perceived for successful completion of the project and the development of an operational multilevel secure local area network. It is understood that some of the specified requirements are designated as mandatory to fulfill near-term functionality and are to be addressed in the initial design. Other requirements, that are annotated, are considered to be future goals and are recorded to support long-range design specifications. This requirements document is intended to provide sufficient detail and content to assist the design team in specification definition.

A.2 System Overview

A.2.1 MLS LAN System Overview

The MLS LAN Project is an effort to provide government and commercial organizations with a cost effective, multilevel, easy-to-use office environment leveraging existing high assurance technology [Ref. 1]. The goals of the project are to produce a networking environment that provides concurrent high assurance access for network users to data at multiple sensitivity levels through the incorporation of inexpensive commercial personal computers.

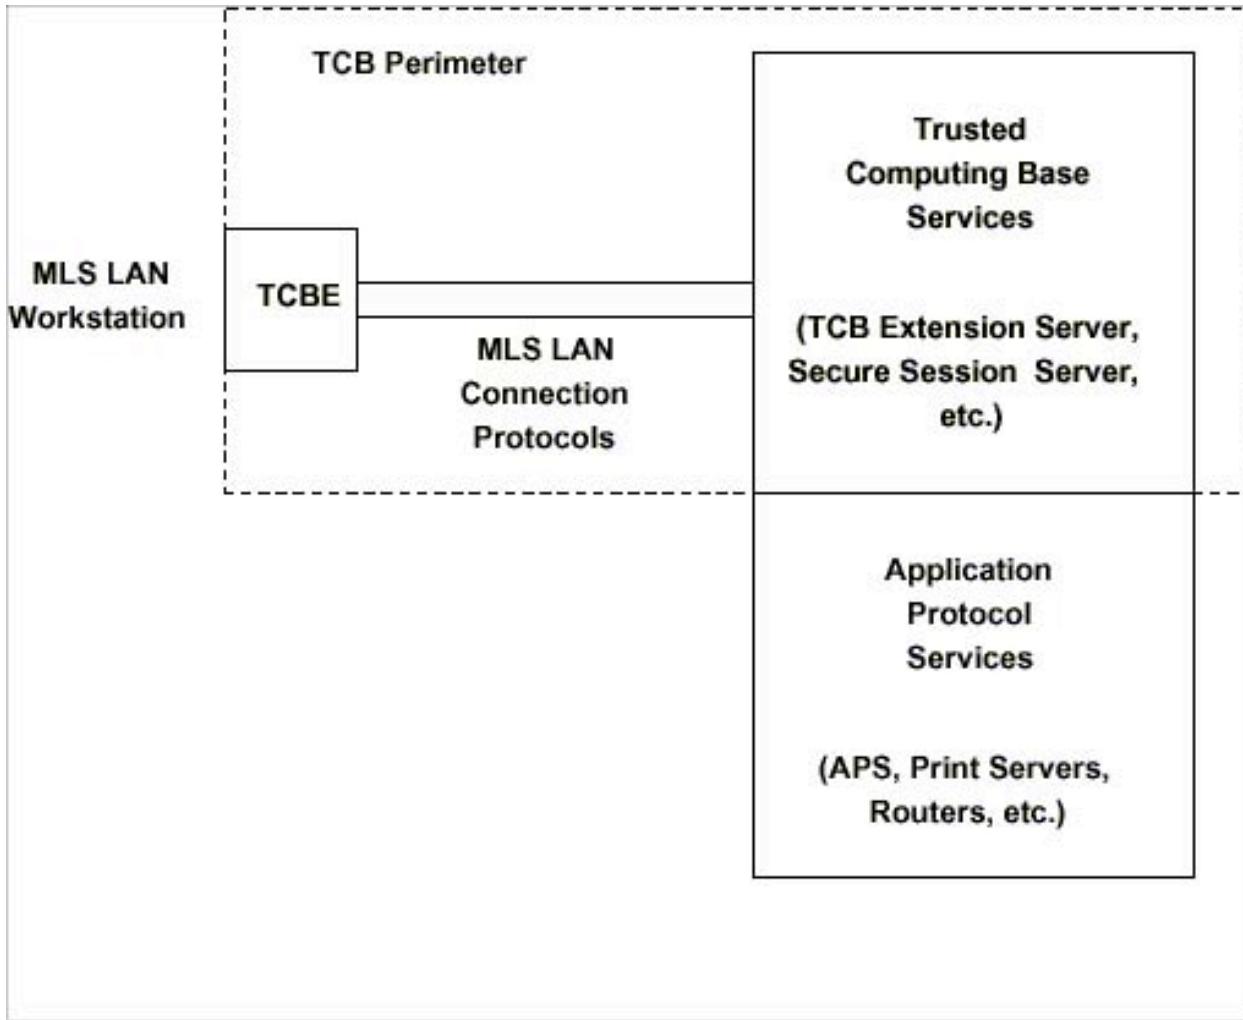


Figure A.2.1. MLS LAN Component Overview

The proposed systems architecture for the MLS LAN is based on the use of the Wang Government Services Incorporated XTS-300(tm) Class B3 rated server. [Ref. 2] The XTS-300 provides both mandatory and discretionary access controls, which "allow separation of users who are at different clearance levels, and prevents a lower level user from reading a higher level user's files or data". [Ref. 3] In accordance with the TCSEC Class B3 rating requirements, the XTS-300 establishes a "Trusted Computing Base" (TCB) that contains all of the Trusted Software Commands, the TCB System Services (TSS), and the Security Kernel. It is the last that implements the TCSEC defined Reference Monitor concept in the XTS-300 [Ref 4]. The MLS LAN incorporates a "logically isolated and unmistakably distinguishable" trusted communications path between the server and its clients through development of a Trusted Computing Base Extension (TCBE). The TCBE will provide a trusted network interface entity for verifiable expansion of the TCB over the communications path to the client workstation. The current hardware solution for the TCBE is to be developed using the Intel I960jx or comparable processor. The TCBE will dominate all actions of the untrusted workstation and allow connectivity into the High Assurance LAN only following the establishment of a trusted path.

A.2.2 MLS LAN User Description

The MLS LAN user is any operator, regardless of authentication, who accesses MLS LAN resources or network functionality. A TCB Authenticated user is one who has successfully established a TCB-to-User connection and been validated by the TCB for operations within the MLS LAN. A Non-TCB Authenticated User, which is a future requirement, is one who has not been validated by the TCB. Accountability of Non-TCB Authenticated Users shall be provided using existing commercial authentication and identification mechanisms.

A.2.3 Component Descriptions

The MLS LAN is comprised of three components (Fig 2.1). The principal component is the Trusted Computing Base (TCB), which provides a fixed security perimeter for MLS LAN operations. Network functionality for access to available application software, file transfer, electronic mail, or remote printing is provided by the Network Application Protocol Services. Finally, the MLS LAN requires a workstation that acts as an agent for the User to access any required network functionality.

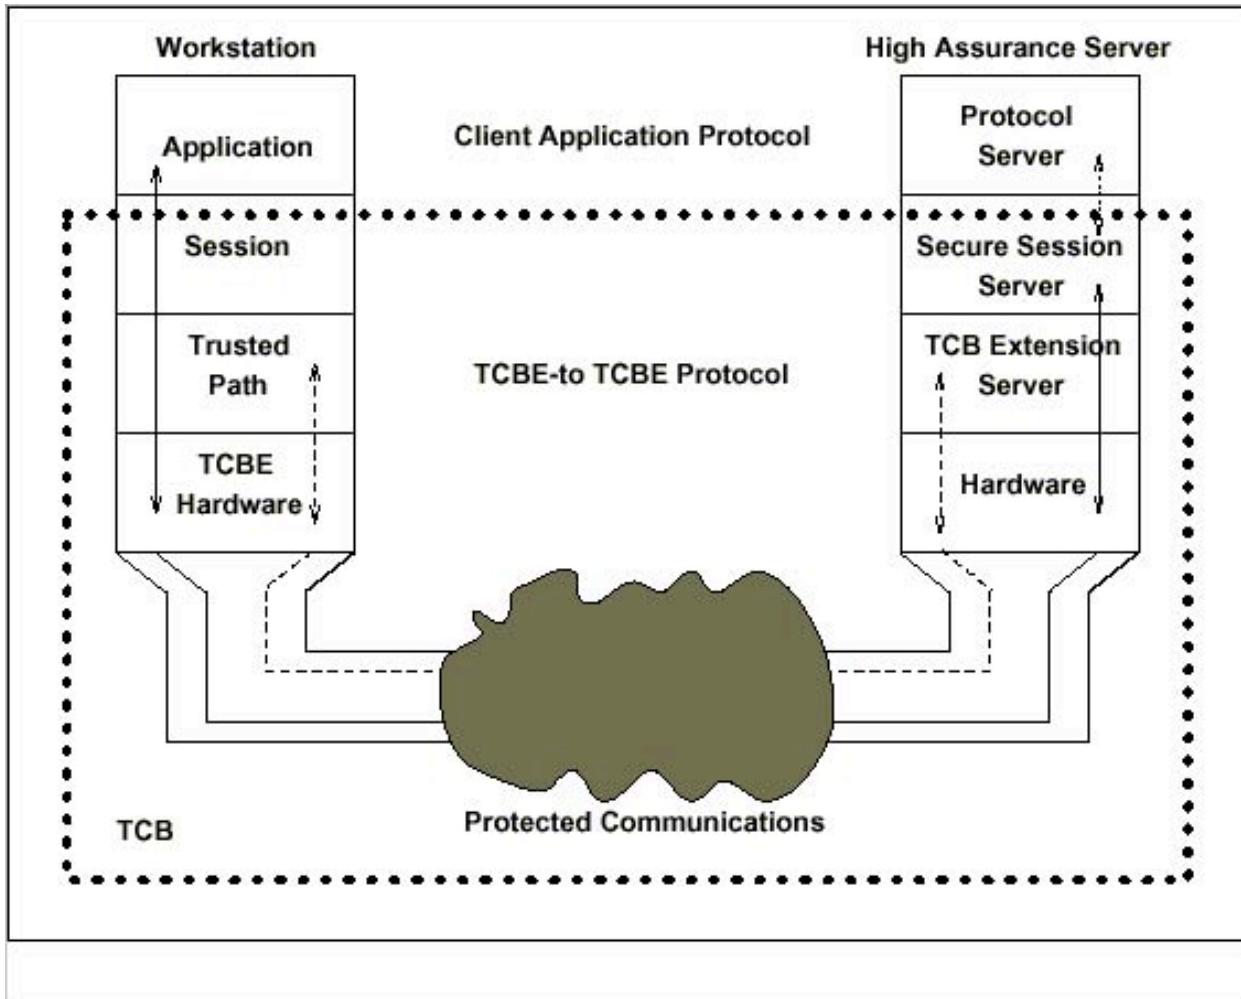


Figure A.3.1. TCB Layering Abstractions

A.2.3.1 Trusted Computing Base

The Trusted Computing Base is an abstraction for the collection of elements of a computer system that pertain to the security policy. Its aegis encompasses all policy enforcement mechanisms, any auditing (retrieval and analysis), identification and authentication, and the interface for security administration.

A.2.3.1.1 Trusted Computing Base Services

The services provided by the MLS LAN to establish a Class B3 rated Trusted Computing Base were outlined in section 2.1 "MLS LAN System Overview". To extend this TCB securely to users additional services are required.

A.2.3.1.2 TCBE Extension Server

The use of the XTS-300 High Assurance Server enables the MLS LAN to place a trusted daemon process in the Operating System Services (OSS) Domain that can provide the protection and communications protocols necessary to establish a trusted path between the workstation and MLS LAN. This "Server" process is used to extend the TCB perimeter securely over the network to the requesting TCBE-equipped workstation. This "Server" process will provide the following functionality: user identification and authentication, session negotiation, session activation, and session termination. [Ref 5.]

A.2.3.1.3 Secure Session Server

The Secure Session Server is an additional trusted daemon "Server" process contained in the OSS. This process will only accept incoming Network Application Protocol Service requests from workstations/users that have established a session via the trusted path and the TCBE Extension Server. Validated requests will be passed on to untrusted Application Protocol Servers, operating on behalf of the user, at the user's negotiated session sensitivity level [Ref 5.]

(Future Requirement) The Secure Session Server will accept Network Application Protocol Services requests from workstation/users that have not established a session, viz. Non-TCB Authenticated Users. These requests will be passed on to untrusted Application Protocol Servers, operating as a system defined anonymous user, at a system defined low secrecy, low integrity, session sensitivity level.

A.2.3.1.4 MLS LAN Session Database Server

The MLS LAN requires a trusted database to maintain all pertinent information concerning each unique TCB session connection. The Session Database Server must provide protection for trusted "read" functionality from all TCB entities and "write" functionality from the TCBE Extension Server.

A.2.3.1.5 Trusted Computing Base Extension

The Trusted Computing Base Extension (TCBE) is a hardware-based computer subsystem that is embedded into the MLS LAN workstation. The TCBE provides the MLS LAN with a verifiable high assurance entity that can be used to extend the TCB.

A.2.3.1.6 MLS LAN Connection Protocols

The MLS LAN connection protocols define the parameters for initiation, security and communications establishment between two or more components of the MLS LAN.

A.2.3.2 Network Application Protocol Services

The MLS LAN uses the TCP/IP stack to support numerous Application Layer Protocol services such as Hyper Text Transfer Protocol (HTTP), Internet Message Access Protocol (IMAP), and File Transfer Protocol (FTP). These services are provided to the users through Application Protocol Servers (APS). While use of these application services are considered "untrusted" and external to the TCB, their access is controlled strictly through the Secure Session Server allowing access to data of multiple sensitivity levels.

A.2.4 MLS LAN

The MLS LAN workstations are the network computers employed by the user to access MLS LAN resources and network functionality.

A.3 System Requirements

A.3.1 MLS LAN Requirements

A.3.1.1

The MLS LAN shall support multiple simultaneous workstation connections.

A.3.1.2

The MLS LAN shall support simultaneous high assurance access for unique workstations operating at different sensitivity levels.

A.3.1.3

The MLS LAN shall provide access to shared resources, application protocol services, and popular application products for both TCB Authenticated Users and, in the future, Non-TCB Authenticated Users.

A.3.1.4

The MLS LAN shall provide high assurance connectivity to application protocols that give access to multiple levels of data in accordance with security policies.

A.3.2 Trusted Computing Base Requirements

This section elaborates on the requirements for the TCB in total. The overall requirements are germane to each of the sub-components while their specific requirements are contained in subsequent sections. An abstract depiction of the MLS LAN layering is provided in Figure 3.1.

A.3.2.1 TCB Overall Requirements

A.3.2.1.1

The TCB shall provide a Secure Attention Key (SAK) mechanism to invoke a trusted path from workstations to which the TCB has been extended.

A.3.2.1.2

The TCB shall establish a trusted path communications connection between network users and the Trusted Computing Base. This trusted path shall be established for initial session authentication purposes, such as "login" or for any specified user operations that require a trusted path, such as "logout", "set session level", downgrade, change user password, etc.

A.3.2.1.3

Once the session has been established, the TCB shall not allow the TCB-to-TCBE Protocol Channel to be broken without loss of network functionality with respect to shared resources, protocol services and applications provided by the MLS LAN.

A.3.2.1.4

The TCB shall allow the user to change the current session sensitivity-level up to the configured maximum for that user.

A.3.2.1.5

The TCB shall provide assurance that the security policy will be enforced in the presence of malicious software.

A.3.2.1.6

The TCB shall provide protection against disclosure and modification of information on all communications channels used by the network.

A.3.2.1.7

The TCB shall control access all devices and networks external to the MLS LAN.

A.3.2.1.8

(Future Requirement) The TCB shall limit the allowable session sensitivity-level to the greatest lower bound between the user's clearance and the TCBE security rating.

A.3.2.2 *Trusted Computing Base Extension Requirements*

A.3.2.2.1

The TCBE shall support the use of Trusted Path communications with the TCB for security related operations.

A.3.2.2.2

The TCBE shall prevent data retention between session security levels and support proper object reuse.

A.3.2.2.3

The TCBE shall support a hardware mechanism that has the ability to purge all memory between session security levels.

A.3.2.2.4

The TCBE shall maintain the ability to reset the host computer system.

A.3.2.2.5

The TCBE shall support the use of a secure attention key.

A.3.2.2.6

The TCBE shall control the information flow into and out of the host computer system.

A.3.2.3 *MLS LAN Connection Protocol Requirements*

A.3.2.3.1

The MLS LAN shall provide a protocol that supports both the establishment of a secure interaction communications channel and the mutual authentication between two TCB entities. This protocol will be known as the "Protected Communications Channel (PCC) Protocol". This protocol will establish the security conduit through which all other MLS LAN protocols operate.

A.3.2.3.2

The MLS LAN shall provide a protocol to support communications between a TCBE equipped workstation and the TCB Extension Server. This protocol will be known as the "TCB-to-TCBE Protocol".

A.3.2.3.3

The MLS LAN shall provide a protocol to support the secure transfer of information from the TCB Extension Server to the Session Database Server to initialize or modify the data maintained on each User Session. This protocol will additionally support the query by a TCB Entity to the Session Database Server for information concerning a User Session. This protocol will be known as the Session Status Protocol.

A.3.2.3.4

The MLS LAN shall provide a protocol to support a TCBE equipped workstation connection to a MLS LAN Secure Session Server. This protocol is the conduit for application protocols and will be known as the "TCBE-to-Session Server Protocol".

A.3.2.3.5

(Future Requirement) The MLS LAN shall provide a protocol to support the connection of a workstation that is not using TCBE services to an untrusted Application Protocol Server, e.g., INTERNET or WWW.

A.3.2.3.6

(Future Requirement) The MLS LAN shall provide a protocol to support a connection of a workstation that is not using TCBE services to a MLS LAN Application Protocol Server.

A.3.3 *MLS LAN Network Application Protocol Services Requirements*

A.3.3.1

The MLS LAN shall support multiple simultaneous accesses to higher layer application protocols, e.g., HTTP, IMAP or FTP.

A.3.3.2

The MLS LAN Application Protocol Servers shall provide access to shared network resources, and popular application products for TCB authenticated users.

A.3.3.3

Access to data maintained on the MLS LAN Applications Protocol Servers (APS) shall be controlled by the TCB in accordance with the security policy.

A.3.3.4

(Future Requirement) The MLS LAN Application Protocol Servers shall provide access to shared network resources, and popular application products for Non-TCB authenticated users.

A.3.4 MLS LAN Workstation Requirements

A.3.4.1

The MLS LAN shall support the use of two configurations of inexpensive commercial personal computers:

A.3.4.1.1

Trusted Computing Base Extension (TCBE) equipped.

A.3.4.1.2

(Future Requirement) Non-TCBE equipped.

A.3.5

The MLS LAN Workstations shall support up-to-date commercial operating systems.

A.3.6

The MLS LAN TCBE Equipped Workstation shall be [in effect] "diskless thin-client" computers operating under the control of the TCBE.

A.4 MLS LAN System Restrictions

A.4.1 MLS LAN Restrictions

A.4.1.1

The MLS LAN shall support no more than one logged in user per workstation at a time.

A.4.2 Environmental Restrictions

A.4.2.1

The TCB platform shall not be subjected to physical tampering.

A.4.2.2

The TCBE hardware shall not be subjected to physical tampering.

A.5 Abbreviations, Acronyms, and Definitions

Abbreviations, Acronyms

APS - Application Protocol Server

CISR - Center for INFOSEC Studies and Research

FTP - File Transfer Protocol

HTTP - Hyper Text Transfer Protocol

IMAP - Internet Message Access Protocol

LAN - Local Area Network

MLS - Multilevel Secure

NPS - Naval Postgraduate School
OSS - Operating System Services
SAK - Secure Attention Key
TCB - Trusted Computing Base
TCBE - Trusted Computing Base Extension
TIC - Trusted Interaction Channel
TSS - TCB System Services
TCSEC - Trusted Computer Security Evaluation Criteria

Definitions

Trusted Computing Base: The Trusted Computing Base is defined as "The totality of protection mechanisms within a computer system - including hardware, firmware, and software - the combination of which is responsible for enforcing a security policy. A TCB consists of one or more components that together enforce a unified security policy over a product or system. The ability of a trusted computing base to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user's clearance) related to the security policy" [Ref 4.]

Trusted Path: The Trusted Path is defined as "A mechanism by which a person at a terminal can communicate directly with the Trusted Computing Base. This mechanism can only be activated by the person or the Trusted Computing Base and cannot be imitated by untrusted software." [Ref 4.]

Session: A Session is defined as the period of interaction between a user and entities within the MLS LAN following session activation and until session termination. Sessions are established or denied based upon based on "attributes such as the location or port or access, the user's security attribute (e.g., identity, clearance level, integrity level, membership in a role), ranges of time (e.g., time-of-day, day-of-week, calendar dates) or combinations of parameters." Limitations may be placed upon user active sessions such as limitations of the number of multiple concurrent sessions or session locking based upon inactivity. [Ref 6.]

TCB Authenticated User: A TCB Authenticated user is one who has successfully established a TCB-to-User connection and been validated by the TCB for operations within the MLS LAN.

Non-TCB Authenticated User: A Non-TCB Authenticated user is one who has not been validated by the TCB.

A. References

1. Irvine, C. E., Anderson, J. P., Robb, D. A., and Hackerson, J. High Assurance Multilevel Services for Off-The-Shelf Workstation Applications. In Proceedings of the 20th National Information Systems Security Conference, pp. 421-431, Crystal City, VA, October 1998.
2. Downey, J., Robb, D. Design of a High Assurance Multilevel Secure Mail Server (HAMMS), Naval Postgraduate School, Monterey CA, September 1997.
3. XTS-300, STOP 4.4.2, Trusted Facility Manual, Document ID: FS96-371-07, Wang Government Services Inc.
4. Department of Defense Trusted Computer System Evaluation Criteria, DoD 5200.28-STD, National Computer Security Center, December 1985.
5. Bryer-Joyner, S., Heller, S. Secure Local Area Network Services for a High Assurance Multilevel Network, Naval Postgraduate School, Monterey, CA. March
6. 1999Common Criteria for Information Technology Security Evaluation Version 2.1, Common Criteria Project Sponsoring Organisations, August, 1999